
Case Study

Brian Kelly

handles EVP products and marketing at KANA.

Greg Bowman

is EMEA alliances manager at KANA.

Service resolution management — A 'knowledge-powered' approach

Brian Kelly and Greg Bowman

Received: 12 May 2005

Abstract

This is a study of the results of applying an interesting approach — embodied here within a specific software suite — to service resolution management within a client's contact centres.

TD Waterhouse UK, part of the Toronto Dominion Bank Group, is one of the UK's largest execution-only brokerages. The business has 3.3 million customer accounts and assets under management of \$130bn, and handles daily trading volumes in excess of \$15m. TD Waterhouse UK also operates share-trading services for customers of NatWest and the Royal Bank of Scotland. The firm enables customers to trade stocks and shares through exchanges across the globe.

To improve the operating efficiency of its two contact centres, TD Waterhouse implemented KANA IQ, a knowledge-base application that lies at the heart of KANA's service resolution management (SRM) suite. With KANA IQ in place throughout its customer support organisation, the firm has a truly multiskilled customer service workforce. Every agent can resolve customer service and sales calls across the UK client's three lines of business — quickly and efficiently.

TD Waterhouse planned to go live with KANA IQ in February 2004. The countdown included a vigorous internal communications campaign to secure buy-in to the project, codenamed 'Project Resolve', from all levels of the business and to gather content for the new knowledge base. TD Waterhouse's objective was to use the IQ knowledge base across three brands as the *only* source of information for its customer service agents. That meant heavyweight input from every specialist in the firm.

The results speak for themselves. Just eight months on, the knowledge base is being used to resolve successfully over one-third of all calls. Average call length has fallen to under three minutes. Best of all, importantly, customer approval rose 16 points in just three months.

Keywords: CRM, SRM, knowledge management, multichannel, e-service, service, customer care, call centre

Greg Bowman

KANA

KANA House

Third Avenue

Marlow, Bucks SL7 1EY, UK

Tel: +44 7818 062 125

Fax: +44 (0)1628 499301

E-mail: gbowman@kana.com

About TD Waterhouse

TD Waterhouse UK operates two contact centres — located in Leeds and Scotland — with a total of 150 seats. Together, they take 175,000 calls each month. Most trading is carried out online, but the contact centres support the online trading platforms and handle more complex queries related to investing in markets outside the UK.

Following the dot.com bust of 2001 onwards, the market for share-dealing services became increasingly competitive and unpredictable. TD Waterhouse has therefore moved to support a wider range of services and customers. To enhance its contact centres' business value, it made four changes to its customer support organisation.

Critical to the project was to focus on the employee experience

- Introducing contact centre 'multiskilling', so that representatives handle both sales- and support-related calls.
- Introducing of a flexible/scaleable full-time equivalent (FTE) model in the contact centre.
- Introducing 'white-label' services.
- Moving 'quote' calls to automated interactive voice response (IVR).

The changes paved the way for TD Waterhouse to become the UK's number one brokerage. But they threatened further complexity and stress for reps. Less-experienced staff would have to handle more complicated queries across multiple lines of business and different customer profiles. Compounding these changes was a rapid increase in recruitment to the contact centres. What was needed was a robust, *process-driven* approach to service resolution, to ensure that new team members could deliver optimal customer service despite minimal prior knowledge of TD Waterhouse's business and products.

Critical to Project Resolve was a focus on the *employee* experience. The aim was to increase reps' productivity — and reduce their turnover — by providing highly effective resolution management. The firm selected KANA to provide the technology platform and implement a knowledge solution that would give its reps easy access to customer support information at the lightning speed of the financial markets.

To minimise the firm's exposure to the financial risks of share buybacks, rapid access to information was vital. A fall in the share price between a trade's confirmation and its completion leaves the firm liable to the customer for any loss incurred. The greater the delays in feeding information back to customers, the greater the risk. Since the KANA IQ knowledge base could provide all details requested by customers at the very moment trades are made, TD Waterhouse anticipated its buyback liabilities to be greatly reduced.

Setting the parameters for success

KANA's software is based on mature knowledge-base and web technologies. These have been developed with one aim in mind: improving the customer experience. A key requirement for the solution — given TD Waterhouse's focus on international markets such as the NYSE, Nasdaq, Deutsche Borse and the Paris Bourse as well as UK exchanges — was its ability to help customers comply with the trading and taxation regimes in force in markets around the world.

With the knowledge base in place, TD Waterhouse needed the ability to offer unparalleled telephone support to its growing investor base. Each rep needed to have at their fingertips the information to resolve any type of call across all three lines of UK business.

A key requirement was the ability to help customers comply with trading and taxation regimes . . .

Key metrics included first call resolution rates, wait times, average call handling time, number of calls, escalations, and productivity

When KANA and implementation partner IBM took up the challenge, several key success criteria were established. Customer satisfaction metrics were set. Prior to the KANA implementation, approval rates from TD Waterhouse’s own customer satisfaction surveys (covering reps’ ability to resolve queries and overall telephone demeanour) lay at 67 per cent. That had to improve. ‘First-call’ resolution rates also had to increase, while ‘wait times’ had to be slashed.

TD Waterhouse also needed to be able to measure the impact of the new knowledge base on the contact centres’ business performance. Key metrics included average call-handling times and numbers of calls ‘escalated’ to back-office departments. But above all, the project’s success would depend on improved productivity in the contact centres — on providing service resolution tools to answer more calls, more successfully, more quickly.

‘Doing the knowledge’

During February and March 2004, TD Waterhouse implemented KANA IQ with the aim of commencing a four-week trial on its Royal Bank of Scotland business line by June. The implementation included three elements:

- populating the knowledge base with the requisite content to help reps resolve customer queries
- securing buy-in to the new system from existing reps
- developing training methodologies to educate and build confidence among new recruits.

Strict targets were set for the creation and updating of new knowledge-base content. Since this would become the source information reps would draw on while speaking to customers, it had to be easy to navigate, clearly written and, above all, authoritative. The three-man TD Waterhouse project team set themselves a goal of creating 1,000 ‘articles’ in seven weeks, with content drawn from experts across the business. Spanning topics from international share trades to transferring ownership of shares following death or marriage, the new content was created, edited and published within KANA IQ’s own business process management tools.

First, the TD Waterhouse team categorised the new content by line of business, then according to reps’ day-to-day needs — such as the most frequently asked questions, or queries that typically would be referred to a supervisor. Once identified, the answers to these questions were written in a familiar style, reducing the risk of an agent having to ask for interpretation.

Once written, content was referred to a ‘virtual team’ of topic experts across the business for approval. This process involved liaising with senior staff from departments including dealing, servicing, back-office operations, compliance, brand, HR, legal and accounts. An audit trail for each sign-off was recorded and ‘ownership’ of content assigned.

The content creation stage was phenomenally successful. Within the seven-week timeframe a staggering 1,700 articles, spanning 40,000 case

The three-man project team set themselves a goal of creating 1,000 articles in seven weeks . . .

Successful . . . a staggering 1,700 articles, spanning 40,000 case types, were published

Those excluded from the trial soon became interested in the solution

types, were created and published. In the constantly changing regulatory environment for global financial services, KANA is enabling new insights on information to be captured, audited and disseminated. Customers always benefit from the insights of the most experienced specialists in the business. As a result, they get their queries resolved quickly and accurately.

While the content team worked on the knowledge base, a substantial internal communications campaign was under way to build TD Waterhouse's customer service reps' awareness of its imminent benefits. This included a 'proof of value' exercise involving a sample of agents. The results and feedback were communicated to the rest of each rep's team weekly. Those excluded from the trial soon became interested in the solution, and the fact that it was reported to be useful and easy to use.

The development team — TD Waterhouse, KANA and IBM — also encouraged potential users' involvement by keeping the project workspace open to all. Agents could see at first hand what was being implemented and understand the reasons why. They also had the chance to navigate around the KANA system, which was customised to reflect the look and feel of TD Waterhouse's business. That provided instant feedback and new ideas.

Four weeks before going live, team communication sessions were held and the solution demonstrated to team leaders and senior managers. Rep training sessions followed, delivered to groups of eight. During the run-in, 'commemorative' stationery, balloons, t-shirts and other paraphernalia stoked anticipation still further.

The results: Knowledge is power

Again, the results of the trial phase were impressive. Overall call-handling times fell 8 per cent, back-office calls fell 20 per cent and the number of customer 'on hold' calls dropped 25 per cent. The pilot also satisfied the service resolution aspect of TD Waterhouse's KANA IQ solution: the ability quickly to provide agents with the answers they needed in 90 per cent of cases.

Following the pilot, the system rolled out to the rest of the contact centre team so as to cover all three of Waterhouse's business lines. The results — which compare the picture before and three months after the implementation — are summarised in Table 1.

The KANA IQ solution was shown to have relieved Waterhouse's back-office departments of significant pressure. All key metrics showed some improvement. The most pronounced reductions were in errors and buybacks caused by, say, incorrectly or slowly executed trades. SRM has delivered TD Waterhouse with major returns and huge savings.

The philosophy of SRM

The economics of SRM are notoriously uneven. There are three key steps of customer service transactions — routing, case tracking and service resolution (Figure 1). But in most contact centre environments there is heavy bias towards the 'mechanics' of call routing and answering. Yet

Call handling times fell 8 per cent, back office calls fell 20 per cent, customers on hold fell 25 per cent

Table 1: Results of the project

	Pre-Resolve	Post-Resolve	
Phone call benefits	Average	Average	% increase/decrease
Average call duration (CC reps)	03:06	02:56	-5.38
Calls made to back-office departments	3,159	2,636	-16.55
Reduction in errors/buybacks			
Number of errors	109	42	-47.96
Number of buybacks	480	374	-22.05

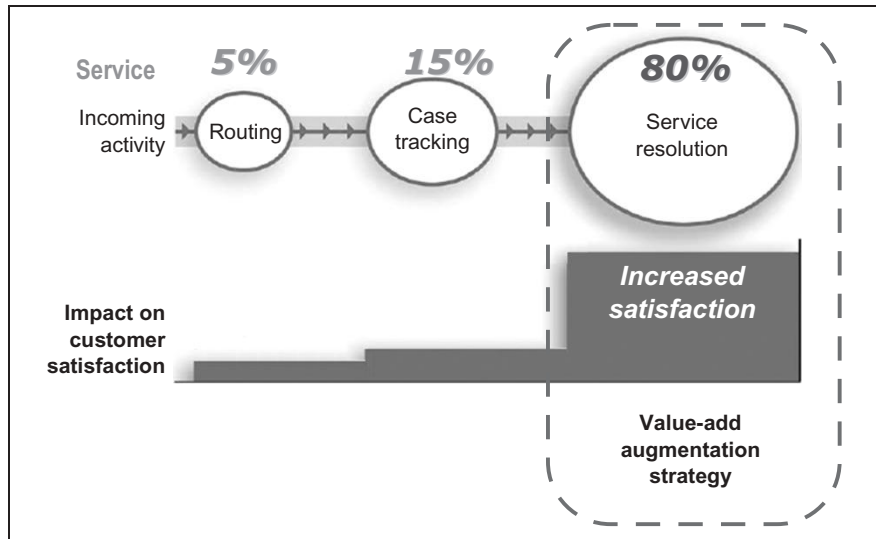


Figure 1: Customer service transactions

Service resolution is critical to the customer experience

routing of calls accounts for only 5 per cent of the cost of providing customer service, according to the SSPA.¹

The service resolution stage (escalating calls, discussing with colleagues, developing responses) is therefore critical to the customer experience. Get that right and businesses are rewarded with the customer's continued loyalty. But a breakdown in the process — whether due to delays, poor internal processes or simply bad management — can prompt customers to move their business elsewhere.

Several factors can cause these breakdowns in service resolution: disparate information; high staff turnover; a lack of repeatable, guiding processes. The statistics are eye-opening, and are causing organisations to take a closer look at the processes that drive their customers' overall service experience. By optimising the resolution process, a business can dramatically cut the overall cost of providing service — and create happier, more loyal customers at the same time.

Routing and case management are mature segments of customer service. Routing is a well-automated process. Most companies already having a system to route service calls properly and efficiently to the right department and agent. Obviously, this is a key component to customer support. But it is only the first step.

The next step is case management, designed to handle customer

The real goal of customer service is to give the customer answers. Call routing and case management provide only part of the solution

contacts or correspondence. Case management tools allow companies to build a record of contacts and customer history. That way, the enterprise can keep up to date and begin to build a relationship with the customer.

Building a successful framework for SRM requires a sound knowledge base of customer data and information. Customer service reps need a knowledge base that gives them a holistic view of the customer life cycle. Information should include how many times a particular customer has called, what type of calls they have made in the past and which products and services they have purchased. These tools are an essential part of the customer service process and are absolutely necessary to automate and manage a call centre. Yet they are only the beginning of the customer experience process. The real goal of customer service is to give the customer answers. Call-routing automation and case management provide only part of the solution.

The third, most vital, step in a customer service transaction is service resolution. To resolve a customer issue or request, customer service agents have to go through five different stages: discovery, research, collaboration, capture and resolve.

- Discovery: understanding the context of the customer issue and the relevant circumstances under which a problem has occurred.
- Research: searching for a solution to a customer problem — often continued as offline research and testing if the issue cannot be determined during the first customer contact.
- Collaboration: joint communication with known expert resources such as discussion threads, tier-two support or formal escalations.
- Capture: the process of recording the data and knowledge developed during the problem resolution process for analytical use.
- Resolution: the final stage of the service resolution process is the creation, delivery and where necessary installation of the customer solution.

Once armed with the correct answer, the agent then needs to deliver the solution to the customer. This can involve providing the ‘best solution’ and offering alternatives, while identifying every possible opportunity to cross-sell or up-sell — though legal reasons may restrict some financial institutions in this final stage. A service resolution application can interface with existing case management and/or call centre applications and optimise the existing process. By understanding the customer’s request and providing the customer service agent with the right tools, information and guidance, it can solve the customer issue.

Complete service resolution involves multiple touch points

A complete service resolution process can involve several touchpoints with the customer — phone, e-mail or self-service, to name but three. The service agent’s job is easier when these communication tools integrate with tools such as search, collaboration, authoring, response and knowledge bases. Workflow processes on top of each solution further streamline the process, reducing time spent.

In summary, a service resolution strategy gives agents quick access to the right information, cutting training costs and significantly reducing

agent turnover. Better service, with quick, consistent and accurate answers, improves customer satisfaction and retention.

The next section will consider the different methods for retrieving the information needed, with special emphasis on the uneasy balance between the need for detailed insights and the low prior knowledge that characterises contact centre environments.

Retrieving the knowledge

The true mark of a knowledge base's potential is its ability to get the right content to the user requesting it. The process is fraught with complications. Among these, lack of knowledge and problems with spellings can lead to particular confusion.

Table 2 summarises the dominant knowledge retrieval techniques, with their advantages and their drawbacks.

Each technique has significant drawbacks. KANA IQ, however, relies on a novel approach to knowledge retrieval: expert reasoning. Expert reasoning makes KANA particularly applicable to customer service scenarios.

Expert reasoning makes use of a series of steps. First, stating the general problem ('My car won't start'). Then, refining the problem definition to make it more specific ('My car won't start because it is out of fuel'). Finally, determining a solution ('I need to fill the tank'). Along the way are questions or tests ('Does the engine turn over?'; 'Does the battery work?'; 'What does the fuel gauge say?'; 'How far have you driven since last refuelling?' etc) to help define the specific problem and the correct solution.

The same steps can be seen at work in any typical customer service interaction. In most cases, the support representative is presented with a general description of the problem. Their job is to ask the customer questions until the specific problem and probable cause are determined. Once they have determined the specific problem, the rep can usually solve it by using their existing knowledge to propose an appropriate solution.

Expert reasoning makes use of a series of steps. The same steps can be seen at work in any typical customer service interaction

Table 2: Knowledge retrieval techniques

Technique	Description	Advantages	Drawbacks
Natural language retrieval	Keyword-based searching	No need to worry about the structure of the knowledge base or the path that a user will take to find the information	Provides no guidance to an end user; relies on some prior knowledge for successful use
Case-based reasoning	Developer describes problem with text and then answers questions so that most likely answers are given greater weight	Easy to build	Can become harder to maintain as more and more cases are added; each question may return a large number of potential answers
Decision trees	User is led to the correct answer by simply answering a series of questions	Easy to use for reps with minimal prior knowledge	'Line of questioning' approach needed can rapidly become tedious for caller and service representative alike
Expert modelling	User provides detailed description of the problem	Highly accurate	Requires the most up-front investment in knowledge-base planning and knowledge acquisition; relies on clearly articulated question from user

The entire problem-resolution process can be described with a few simple terms.

- What is the *problem*?
- What is the *cause*?
- What is the *solution*?
- What *questions* can I ask?

Expert reasoning is built around these four simple objects: problems, causes, solutions and questions. These are the same terms used by every customer service representative. In fact, one of expert reasoning's greatest strengths is that it allows organisations to build a knowledge base while continuing to work with these familiar terms and concepts. All they need do is define a set of problems; a set of solutions; and the questions that will help them link problems with solutions.

As Figure 2 shows, the 'problem' lies at the centre of the service resolution process. The more specific the problem definition, the easier it is to determine both the cause and the solution. Other methodologies talk about cases or branches that lead to a solution, but expert reasoning makes the problem the central object within the knowledge base. This has two advantages: it makes it easier to build a knowledge base; and it makes it easier for an end user to find the information they need.

A user can begin an expert reasoning session by answering an open-ended question with a string of text, or by answering a series of questions (Figure 3). Either method helps define the general problem. They will then be presented with a series of questions — which may combine aspects of case-based reasoning, decision trees and expert models — to refine the general problem quickly into a very specific problem (the inner arrows in Figure 3). In each case, the expert reasoning algorithm automatically selects the best set of questions to narrow the problem definition efficiently.

Once the system has determined a specific problem definition, it can usually provide an answer very quickly. In some cases, the specific problem definition is sufficient to lead directly to a solution. In other cases, the system needs to employ a further set of questions to determine the most appropriate solution for the end user's problem. Once again, the expert reasoning algorithm chooses the most efficient method to find the correct answer as quickly as possible.

A user can begin by answering a series of questions to define the general problem




Figure 2: The problem-resolution space employed by expert reasoning

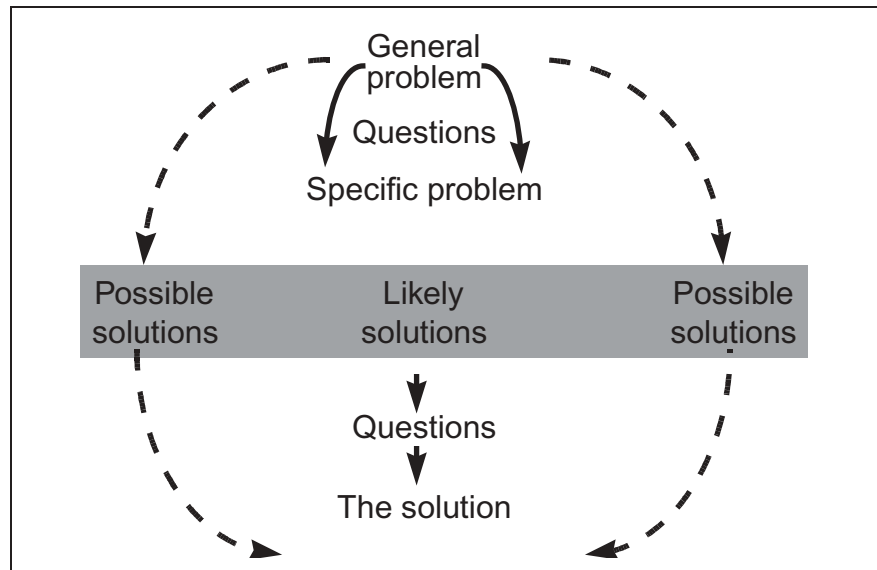


Figure 3: The possible paths to a solution with expert reasoning

As a result, a customer service representative can deliver expert knowledge to customers despite having next to no awareness of their employer's business, products or customers. In fact, they only really need to understand the service resolution *process*, which guides them from the customer's original question to the solution they actually need.

Conclusion

TD Waterhouse's experiences with SRM highlight how a highly targeted investment in a single part of the customer service process can pave the way for a wide range of high-impact business strategies. In this case, the client diversified its business lines and introduced multi-skilled contact centre staff.

This paper has demonstrated the importance of underpinning the strategy with the right technology. SRM can help solve customer service queries of any complexity, with minimal assumed knowledge on the part of the rep handling the call. But the information required has to be authoritative, up to date and — above all — easily retrievable in real time.

Reference

1. Service and Support Professionals Association, available at www.thesspa.com.

TD Waterhouse diversified its business lines and introduced multi-skilled contact centre staff